

Plurality Voting under Uncertainty

Reshef Meir

November 19, 2014

Abstract

Understanding the nature of strategic voting is the holy grail of social choice theory, where game-theory, social science and recently computational approaches are all applied in order to model the incentives and behavior of voters.

In a recent paper, Meir et al. [15] made another step in this direction, by suggesting a behavioral game-theoretic model for voters under uncertainty. For a specific variation of best-response heuristics, they proved initial existence and convergence results in the Plurality voting system.

In this paper, we extend the model in multiple directions, considering voters with different uncertainty levels, simultaneous strategic decisions, and a more permissive notion of best-response. We prove that a voting equilibrium exists even in the most general case. Further, any society voting in an iterative setting is guaranteed to converge.

We also analyze an alternative behavior where voters try to minimize their worst-case regret. We show that the two behaviors coincide in the simple setting of Meir et al., but not in the general case.

1 Introduction

Suppose that your favorite candidate in the elections trails behind in the polls. A game-theorist advice in such a case would be to “vote strategically.”

As a simple example, consider an election with 3 candidates $M = \{a, b, c\}$. Suppose that the most recent poll indicates that a is leading with 45% of the votes, whereas there are 40% for b and 15% for c (see Fig. 1). A voter who prefers c over b over a may conclude that c has no chance of winning, and thus a wise strategic decision would be to vote for b . However, there is no consensus whatsoever regarding how to generalize this observation to an arbitrary situation, or even on the conditions under which c should be deserted.

This lack of a conclusive notion for strategic voting is partly due to the fact that there are many ways to describe the information voters have, as well as their beliefs over other voters’ preferences and actions when casting their vote. Moreover, even for a given belief there may be several different actions that can be justified as rational. Hence any model of voting behavior and equilibrium must state explicitly its epistemic and behavioral assumptions.

We can of course apply some “standard” assumptions from game-theory, for example that voters play a Nash equilibrium of the game induced by their preferences. However, such a prediction turns out to be very uninformative: a single voter can rarely affect the outcome, and thus almost any way the voters vote is a Nash equilibrium (including, for example, when all voters vote for their least preferred candidate).

It is therefore natural to consider the role of *uncertainty* in voters’ decisions. One of the most prominent models for voting under uncertainty was suggested by Myerson and Weber [19], where voters preferences are sampled from a known prior distribution. An equilibrium according to the MW model is a mapping from preferences to a distribution over votes, such that each voter maximizes her expected utility w.r.t. this distribution. Myerson and Weber prove via a fixed-point argument that an equilibrium always exists for every positional scoring rule, and in particular for Plurality.

The MW model reflects a classical economic approach (as do many other game-theoretic models of voting), but its epistemic and behavioral assumptions seem highly unreasonable in the context of voting. As studies in behavioral psychology show, human decision makers often ignore probabilistic information even when it is given, employing various heuristics instead [23]. It is thus unlikely that people are able to represent such complex distributions, or to compute and optimize their own expected utility, let alone the equilibria of the game.

In a recent effort to reconcile well-founded decision making approaches with a formal game-theoretic analysis, Meir et al. [15] suggested a model for Plurality voting relying on *strict uncertainty* and *local dominance*. Informally, voters’ beliefs can be described by a single vector of candidates’ scores. This prospective score vector may be the result of a poll, derived from acquaintance with the other voters, from the outcome of a previous round of voting, etc. Each voter has a single uncertainty parameter, reflecting how sure she is in the correctness of the prospective scores—higher uncertainty means the voter considers a larger range of outcomes (score vectors) as “possible”. In the example introduced above, a voter with an intermediate uncertainty parameter will consider a tie between a and b as “possible,” but will reject any outcome where c can win as impossible.

Given this uncertain, non-probabilistic view, it is not a-priori clear how a voter should act. Meir et al. adopted one of the classic approaches to decision making under strict uncertainty, assuming that a voter i will refrain from voting for a candidate a_i that is *locally dominated* by another candidate a'_i [1, 2]. Intuitively, a'_i *locally-dominates* a_i if it is always at least as good, and sometimes strictly better, to vote for a'_i than for a_i (taking into account all score vectors that voter i considers possible). A *voting equilibrium* is simply a state where no voter votes for a locally-dominated candidate according to her beliefs. Going back to our example, our voter will consider candidate c as locally-dominated by candidate b , and will thus strategize by voting for b .

Informally, Meir et al. [15] proved that if all voters have the *same uncertainty parameter*, start from the truthful state, and play one at a time, then they always converge to a voting equilibrium. In the special case of zero uncertainty, this

simply means that the best-response dynamics converges to a Nash equilibrium (which is known from [17]).

However, these assumptions are rather restrictive. Some voters may be less informed than others, or simply more “stubborn” (e.g., require even fewer votes to c in order to be convinced it cannot win). Also, there is no reason to believe that voters initially vote for their most preferred candidate, or that only one voter changes her vote after each poll or voting round (see, e.g., [20]). Meir et al. [15] showed *empirically* that convergence occurs even without those assumptions, and conjectured that at least some of the assumptions could be relaxed. However they provide no formal guarantee for such convergence, and in fact for voters with different uncertainty levels even existence of equilibrium remains an open question.

1.1 Research goals

The main purpose of this work is to close this gap, and to prove equilibrium existence and convergence under conditions that are as broad as possible.

While we adopt the model of Meir et al. at large, we study a non-atomic variation of it, where the effect of any single voter is negligible (as in [19]). This simplifies the model, and allows us to prove stronger results when arbitrary subsets of voters change their vote simultaneously. However all of our results hold in the finite case, for voters that move one-at-a-time (see Appendix VIII). We modify the distance function used to determine the possible ranges of candidates’ scores, to one that is better justified by psychological studies (for details see Footnote 3).

In addition, we are interested whether the behavioral assumptions can be weakened, or, alternatively, be replaced with a more nuanced way to select among several undominated candidates. To that end, we define a variation of the model where voters are aiming to minimize their *worst-case regret* over all states they believe possible, and compare this behavior to voting under local-dominance.

1.2 Our contribution

Our main result is proving that voters with local-dominance behavior always converge to an equilibrium. This holds for any mixture of types (including voters with different uncertainty levels), for any initial voting profile, and for any order of moves (including moves of arbitrary subsets of voters, and suboptimal moves).

We then turn to study voters minimizing worst-case regret. When voters have the same uncertainty level and start from the truthful profile, we show that regret minimization coincides with local dominance. However if these requirements are relaxed then the behaviors may significantly differ, and even the existence of equilibrium is not guaranteed.

2 The Formal Model

Basic notations Where possible, we follow the notations and definitions of [15]. We make the necessary changes to adapt their model to nonatomic voters. We denote $[x] = \{1, 2, \dots, x\}$ for all $x \in \mathbb{N}$. For a finite set A , $\Delta(A)$ denotes the set of all probability distributions over A (all non-negative vectors of size $|A|$ that sum to 1).

The set of candidates is denoted by M , where $m = |M|$. Let $\mathcal{Q} = \pi(M)$ be the set of all strict orders (permutations) over M . We encode all the information on a voter, including her preferences, in her *type*,¹ where the set of all types is denoted by \mathcal{V} .

Thus a voter of type $v \in \mathcal{V}$ has preferences $Q_v \in \mathcal{Q}$, where $Q_v(a) \in [m]$ is the rank of candidate $a \in M$ (lower is better), and $q_v = Q_v^{-1}(1)$ is her most-preferred candidate. We denote $a \succ_v b$ if $Q_v(a) < Q_v(b)$.

We do not have a finite set of voters. Rather, a *preference profile* $\mathbf{Q} \in \Delta(\mathcal{Q})$ is a distribution over preferences, specifying the fraction of voters with each preference order. A *population* $\mathbf{V} \in \Delta(\mathcal{V})$ is a distribution over *types*, that is, a preference profile aggregated with any additional information specified by voters' types (e.g., their beliefs and behaviors).

Under the Plurality rule, every voter selects a single candidate. Intuitively, we should specify how many voters of each type vote for every candidate. Formally, an *action profile* (also called *state*) \mathbf{a} is a refinement of \mathbf{V} , where $a(v, c) \in \mathbb{R}_+$ denotes the fraction of voters of type v who vote for $c \in M$. Our next definitions are intended to enable notations and analysis that are similar to those used in the atomic case.

Since a single voter has negligible influence, we only consider moves by subsets of voters whose size is a multiple of ϵ (for some arbitrarily small ϵ). All voters in each set have the same type and move simultaneously (although they need not be aware of this, see also Appendix IX). We denote by I the collection of these $1/\epsilon$ sets. Since all voters in set $i \in I$ are indistinguishable we refer to “voter i ” which is an arbitrary voter in the set i . This voter has a well-defined type v_i and a well defined action a_i in every action profile \mathbf{a} .

Winner determination and tie-breaking The outcome of the Plurality rule in state \mathbf{a} , denoted $f(\mathbf{a}) \in M$, only depends on the total number of votes for each candidate. We thus define the *score vector* $\mathbf{s}_{\mathbf{a}}$ induced by action profile \mathbf{a} . That is, $s_{\mathbf{a}}(c) = \sum_{v \in \mathcal{V}} a(v, c) = |\{i \in I : a_i = c\}| \epsilon$. We will use \mathbf{a} and $\mathbf{s}_{\mathbf{a}}$ interchangeably, sometimes omitting the subscript \mathbf{a} . We denote by $\mathcal{S} = \mathbb{R}_+^m$ the set of all score vectors. We also sometimes refer to score vectors as states, although \mathbf{s} may not be attained from an actual action profile. E.g., it is possible that $\sum_{c \in M} s(c) \neq 1$. Note that we may only use \mathbf{s} in a context where voters' identifiers are not important.

Note that changing the vote of a single voter in \mathbf{a} does not change \mathbf{a} or $\mathbf{s}_{\mathbf{a}}$, yet we need a way to settle tie-breaking. Every score vector \mathbf{s} implicitly contains

¹We later extend the definition of a type to also include the belief structure and behavior of the voter.

an arbitrary “tie-breaker” $Q_{\mathbf{s}}^+ \in \mathcal{Q}$.

The *winner* $f(\mathbf{s})$ is the candidate $c \in M$ whose score $s(c)$ is maximal. If there is more than one candidate with maximal score, we break ties according to $Q_{\mathbf{s}}^+$. Formally, $f(\mathbf{s}) = \operatorname{argmin}_{c \in \operatorname{argmax} s(c')} Q_{\mathbf{s}}^+(c)$. Note that the outcome is well defined even for score vectors that are not derived from valid action profiles.

Dynamics and equilibria We describe the behavior of a voter of type $v \in \mathcal{V}$ by a *response function* $g_v : M \times \mathcal{S} \rightarrow 2^M \setminus \emptyset$. That is, a mapping from the current state (only taking into account aggregate scores) and current action, to a subset of actions. Together, Q_v and g_v completely define the type v . For an identified voter $i \in I$, we write $g_i(\mathbf{s})$ instead of $g_{v_i}(a_i, \mathbf{s})$. We also write $g_i(\mathbf{a})$ as a shorthand for $g_i(\mathbf{s}_{\mathbf{a}})$. Intuitively, this means that a voter $i \in I$ may choose any action in $g_i(\mathbf{a})$. In the terms of [11], g_i is a *historyless*, non-deterministic response function.

Definition 1. A *voting equilibrium* for population \mathbf{V} is a state \mathbf{a} , where $g_i(\mathbf{a}) = \{a_i\}$ for all $i \in I$.

For example, if g_i is the *best-response function* ($g_i(\mathbf{a}) = \{c \in M : c \text{ maximizes the utility of } i \text{ in } \mathbf{a}\}$ for all i), then a voting equilibria coincides with the Nash equilibria of the game. We emphasize that classic results on existence of equilibria in nonatomic games (e.g., [22]) do not apply, since even if we assign cardinal utilities to voters, those would be highly discontinuous in the action profile. Note that as in [19], we do not assume that voters with the same type, or even the same identifier, are coordinated.

Even in cases where an equilibrium exists, players may or may not reach one, depending on their initial state and the order in which they play. We are therefore interested in sufficient conditions under which the game is *acyclic*, i.e. there are no finite cycles of rational moves by groups of any size.²

3 Uncertainty and Strategic Voting

The most important part of the model is of course the way we define the behavior, which is the response function g_i . This behavior depends on voters’ preferences, and also on their beliefs about the current state. We assume voters derive their beliefs using a distance-based strict uncertainty model, following [15].

3.1 Distance-based uncertainty

Suppose we have some distance measure for score vectors, denoted by $\delta(\mathbf{s}, \mathbf{s}')$. For any \mathbf{s} , let $S(\mathbf{s}, x) \subseteq \mathcal{S}$ be the set of vectors that are at distance at most x from \mathbf{s} . Formally, $S(\mathbf{s}, x) = \{\mathbf{s}' : \delta(\mathbf{s}, \mathbf{s}') \leq x\}$. A distance function δ is *candidate-wise* if it can be written as $\delta(\mathbf{s}, \mathbf{s}') = \max_{c \in M} \delta'(s_c, s'_c)$ for some monotone

²In finite games, acyclicity and guaranteed convergence imply one another. In nonatomic games the relation is a bit more delicate, see Appendix IX.




Figure 1: On the left figure we see a given state \mathbf{s} , which can be thought of as the result of a poll, or of the current voting round. On the right we depict the set of possible states $S(\mathbf{s}, r)$ for $r = 0.15$. Any \mathbf{s}' is a possible state as long as the score of each candidate is in the marked range. The dashed line marks the threshold above which a candidate is considered a possible winner.

function δ' (meaning that for a fixed s , $\delta'(s, s')$ is nondecreasing in $|s - s'|$). Thus $\mathbf{s}' \in S(\mathbf{s}, x)$ only if the score of every candidate in \mathbf{s}' is sufficiently close to its score in \mathbf{s} , see Fig. 1.

In [15], several metrics have been suggested, including various ℓ_d norms and the earth-mover distance. Among those the multiplicative distance (where $\delta'(s, s') = \max\{s'/s, s/s'\} - 1$) and the ℓ_∞ metric (where $\delta'(s, s') = |s - s'|$) are candidate-wise.

While all of our results apply for any candidate-wise distance, for concreteness we assume throughout the paper the multiplicative distance. Note that the multiplicative distance is independent of the amount of voters, i.e. $\delta(\mathbf{s}, \mathbf{s}') = \delta(\alpha\mathbf{s}, \alpha\mathbf{s}')$ for all $\alpha > 0$. This is consistent with findings on how people perceive uncertainty over numerical values [13, 23].³

We define the *uncertainty parameter* $r_v \in \mathbb{R}_+$ as part the agent's type (will be later used to construct the response function g_v). Given a profile \mathbf{a} , voter $i \in I$ believes that the actual state may be any $\mathbf{s}' \in S(\mathbf{s}_{\mathbf{a}}, r_i)$, where $r_i = r_{v_i}$.

Note that the distance between two score vectors that only differ by the tie-breaker is 0. Thus if several candidates $C \subseteq M$ are tied in $\mathbf{s} = \mathbf{s}_{\mathbf{a}}$ with maximal score, for each $c \in C$ there is a state $\mathbf{s}' \in S(\mathbf{s}, r_i)$ where $f(\mathbf{s}') = c$ (any state where $Q_{\mathbf{s}'}^+(c) = 1$), and this holds for any r_i .

³For example, the studies show that if the average number of girls born daily in a hospital is s , then people believe that the probability that on a given day the number is within $[(1 - r)s, (1 + r)s]$ is fixed and does not depend on s . Kahneman and Tversky highlight that this reasoning stands in contrast to the scientific truth in this case, where the range r is proportional to $1/\sqrt{s}$. In our case we can think of the score of a candidate as the limit of a Poisson variable (rather than Binomial in the hospital example), and thus the range $[s/(1 + r), s(1 + r)]$ is more appropriate.

A voter facing strict uncertainty may use various heuristics when selecting a strategy. In this work we follow two standard approaches: avoiding dominated strategies [2], and minimizing worst-case regret [21, 10].

Voter influence In the nonatomic case, every voter has negligible influence. However, this influence is not perceived by the voter as zero, as otherwise voters would always be indifferent about their actions. We define $f(\mathbf{s}, c)$ to be the outcome that a voter expects under score vector \mathbf{s} , when voting for c . Intuitively, the extra vote c decides the winner if several candidates are tied with maximal score in \mathbf{s} (overriding the default tie-breaker $Q_{\mathbf{s}}^+$), and otherwise has no effect. A more formal definition appears in Appendix X.

3.2 Local dominance

The first approach we consider follows [5, 15], where voting dynamics is determined by local dominance relations.

Consider a particular voter of identifier $i = \langle v_i, a_i \rangle$. Let $S = S(\mathbf{s}, r_i)$. We say that action a_i *S-beats* b_i if there is at least one state $\mathbf{s}' \in S$ s.t. $f(\mathbf{s}', a_i) \succ_i f(\mathbf{s}', b_i)$. That is where i strictly prefers $f(\mathbf{s}', a_i)$ over $f(\mathbf{s}', b_i)$. Action a_i *S-dominates* b_i if (I) a_i *S-beats* b_i ; and (II) b_i does *not* *S-beat* a_i .

We next define the response function that strategic voters apply under the local-dominance behavioral model. Consider a voter i , where $v_i = \langle Q_i, r_i \rangle$.

Definition 2 (Strategic move under local dominance (LD)). Let a_i be the vote of i in profile \mathbf{a} . Let $D \subseteq M$ be the set of candidates that $S(\mathbf{a}, r_i)$ -dominate a_i , and that are not $S(\mathbf{a}, r_i)$ -dominated. If $D = \emptyset$ then $g_i(\mathbf{a}) = \{a_i\}$. Otherwise:

- For a *weak LD voter*, $g_i(\mathbf{a}) = D$.
- For a *strict LD voter*, $g_i(\mathbf{a}) = \operatorname{argmin}_{d \in D} Q_i(d)$.

In other words, a strict LD voter votes the most preferred candidate that locally-dominates her current choice (if such exists), which is almost identical to the definition in [15] for finite populations (see also Def. 5 in Appendix). The definition of a weak LD voter is much more permissive, and does not restrict the voter to select the most preferred candidate in D . A higher value of r_i may either indicate that the voter is less informed, or simply that she requires stronger evidence that a move will be beneficial, a tendency we can interpret as stronger *loss aversion* [12].

As in the model of [15], we get the following immediate observation.

Proposition 1. *Let \mathbf{V} be a population of LD voters. A profile \mathbf{a} is a voting equilibrium iff $\forall i \in I, \exists a'_i \in M$, such that a'_i $S(\mathbf{a}, r_i)$ -dominates a_i . I.e., if no voter votes for a locally dominated candidate.*

4 Convergence with LD Voters

4.1 Strategic responses and possible winners

We say that candidate c is a *possible winner for i in state \mathbf{s}* if there is a possible state where c wins. Formally, $W_i(\mathbf{s}) = \{c \in M : \exists \mathbf{s}' \in S(\mathbf{s}, r_i) \text{ s.t. } f(\mathbf{s}', c) = c\}$. Also denote $W_0(\mathbf{s}) = \{c \in M : f(\mathbf{s}, c) = c\} = W_i(\mathbf{s})$ for $r_i = 0$, and $W_i(\mathbf{s}) = W_i(\mathbf{s}_a)$.

It is easy to see that under the multiplicative distance $c \in W_i(\mathbf{s})$ iff $s(c) \geq (1 + r_i)^{-2} s(f(\mathbf{s}))$ (see dashed line in Fig. 1), and that similar thresholds exist for other metrics (see also Lemma 3 in [16]).

Lemma 2. *Every pair of possible winners are tied in some possible state. Formally, for every $b, c \in W_i(\mathbf{s})$, there is $\mathbf{s}' \in S(\mathbf{s}, r_i)$ s.t. $b, c \in W_0(\mathbf{s}')$.*

Proof. Consider some $b, c \in W_i(\mathbf{s})$. For the multiplicative distance, let the score of all candidates except b, c be $s'(a) = s(a)/(1 + r_i)$, and set $s'(b) = s'(c) = \min\{s(b), s(c)\}(1 + r_i)$. Then $\mathbf{s}' \in S(\mathbf{s}, r_i)$, and $s'(b) = s'(c) \geq s'(a)$ for all $a \in M$. \square

We denote by $a_i \xrightarrow{i} a'_i$ valid local dominance steps where $a'_i \in g_i(\mathbf{a})$ and $a'_i \neq a_i$. The next two lemmas characterize such LD moves.

Lemma 3. *Consider an LD move $a_i \xrightarrow{i} a'_i$. Then either (a) $a_i \notin W_i(\mathbf{s})$; or (b) $a_i \prec_i b$ for all $b \in W_i(\mathbf{s})$; or (c) $r_i = 0$, $\{a_i, a'_i\} \subseteq W_0(\mathbf{s})$ and $a'_i \succ_i a_i$.*

Proof. Suppose that $a_i, b \in W_i(\mathbf{a})$, and $a_i \succ_i b$ (i.e., (a) and (b) are violated). Assume first that $a'_i \notin W_0(\mathbf{s})$. By Lemma 2 there is a state $\mathbf{s}' \in S(\mathbf{s}, r_i)$ where a_i, b have maximal score (possibly with other candidates), strictly above a'_i . W.l.o.g. $Q_{\mathbf{s}'}^+(b) < Q_{\mathbf{s}'}^+(a_i)$, as the tie-breaker does not affect the distance. Thus $f(\mathbf{s}', a_i) = a_i, f(\mathbf{s}', a'_i) = b$. Since $a_i \succ_i b$, then we have that a_i $S(\mathbf{a}, r_i)$ -beats a'_i .

The remaining case is where $a'_i \in W_0(\mathbf{s})$ and $a_i \succ_i a'_i$. Then in the state \mathbf{s}' where a_i, a'_i are tied it is better to vote for a_i . In either case we get that $a'_i \notin D$, which is a contradiction. \square

Lemma 4. *Consider an LD move $a_i \xrightarrow{i} a'_i$. Then*

1. $a'_i \in W_i(\mathbf{s})$, and there is some $c \in W_i(\mathbf{s})$ s.t. $a'_i \succ_i c$.
2. For a strict LD move, $a'_i = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$.
3. If $a_i \notin W_i(\mathbf{s})$, then $|g_i(\mathbf{s})| = 1$ (thus weak and strict LD coincide).

Part 1 is the only thing we need for our later convergence results, and its proof is immediate: Consider $a_i^* = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$. Clearly a_i^* locally dominates any candidate not in $W_i(\mathbf{a})$, thus $a_i^* \in W_i(\mathbf{s})$. If there is only one possible winner, no action dominates any other action, and there are no moves. Part 2 is the nonatomic analog of Lemma 5.1 in [15], however our proof is simpler (see Appendix XI).

The above lemmas show that a strict LD move boils down to a simple heuristic: vote for the most preferred candidate among those whose score is above the threshold. For a weak LD move, any candidate above the threshold except the least preferred can be selected.

4.2 Existence of equilibrium and convergence

Suppose that voters start at some arbitrary state \mathbf{a}^0 , and play repeatedly. We get a sequence of states $\mathbf{a}^0, \mathbf{a}^1, \dots$ where in every iteration some arbitrary subset of voters may change their vote. That is, for any $t \in \mathbb{N}$ and any $i \in I$, either $a_i^{t+1} \in g_i(\mathbf{a}^t)$, or $a_i^{t+1} = a_i^t$. Our primary result states that voters always converge in the nonatomic model (see Appendix VIII for a proof and discussion of the finite case).⁴

Theorem 5. *Any sequence of weak LD moves is finite.*

Proof. We only need to show that no valid sequence may contain a cycle (see Appendix IX). Assume, toward a contradiction, that there is a cyclic path $(\mathbf{a}^t)_{t=0}^T$, and denote by $R \subseteq M$ all candidates that are part of the cycle. Let s^* be the lowest score of any candidate in R during the cycle, w.l.o.g. candidate $a^* \in R$ at time t^* . Thus $s^* = s^{t^*}(a^*) \leq s^t(c)$ for all $c \in R$ for every time $t \leq T$. Consider the next step where some voters join a^* (w.l.o.g. at step t^*), and pick an arbitrary voter $j \in I$ s.t. $a^* = a_j^{t^*+1} \neq a_j^{t^*}$. Thus at step t^* there is a move $a_j \xrightarrow{j} a^*$ where $a_j = a_j^{t^*+1} \in R$.

By Lemma 4, a^* is a possible winner for voter j , i.e., $a^* \in W_j(\mathbf{s}^{t^*})$. Since $s^{t^*}(c) \geq s^{t^*}(a^*)$ for all $c \in R$, we have $R \subseteq W_j(\mathbf{a}^{t^*})$, and in particular $a_j \in W_j(\mathbf{a}^{t^*})$.

By Lemma 3, either a_j is the least-preferred candidate for j in $W_j(\mathbf{s}^{t^*})$ (Cases I and II below), or the third category of the lemma holds. We treat the latter case separately (Case III), so assume a_j is indeed the least-preferred in $W_j(\mathbf{s}^{t^*})$. Since $R \subseteq W_j(\mathbf{a}^{t^*})$, a_j is the least-preferred in R as well.

There must be some step t^{**} in the cycle where a voter of type v_j moves to a_j (w.l.o.g. voter j). So in $\mathbf{s}^{t^{**}}$, a_j is preferred by j to *some* other possible winner z by Lemma 4. Since a_j is the least-preferred in R , and $a_j \succ_j z$, we have that $z \in W_j(\mathbf{s}^{t^{**}}) \setminus R$. Denote the (fixed) score of z by $s(z)$.

Case I: $s(z) \geq s^*$. Consider again step t^* . Since $s(z) \geq s^*$, we have $z \in W_j(\mathbf{s}^{t^*})$. Since a_j is the least preferred possible winner in t^* , we have that $z \succ_j a_j$, which is a contradiction.

Case II: $s(z) < s^*$. Denote $d = a_j^{t^{**}}$, and consider the step $d \xrightarrow{j} a_j$ at time t^{**} . Since $d \in R$ then $s^{t^{**}}(d) \geq s^* > s(z)$, and thus $d \in W_j(\mathbf{s}^{t^{**}})$. By Lemma 3 (category (b) or (c)), we have that $a_j \succ_j d$. This is a contradiction since a_j is the least-preferred in R .

⁴We assume strict preference orders for consistency with most of the social-choice literature. However we note that the theorem applies also under weak preferences. Indeed, introducing indifference in the preference relation only eliminate LD moves, and therefore never creates a cycle.

Case III: The remaining case is when there is $b \in W_j(\mathbf{s}^{t*})$ s.t. $a_j \succ_j b$. Then by Lemma 3, $a^* \in W_0(\mathbf{s}^{t*})$. However since a^* has minimal score, this means that $R \subseteq W_0(\mathbf{s}^{t*})$, i.e., all candidates in the cycle have the same score s^* at time t^* . Then all of R must have the same score at *every time* t , since if the score of some candidates goes up, the score of others must go down below the minimum s^* . This means that all of the moves in the cycle fall under categories (b) or (c) of Lemma 3. Thus voters only vote for more preferred candidates, which contradicts a cycle. \square

One can argue that the uncertainty level of a voter may not remain the same throughout the game. For example, there may be less uncertainty as the game advances. We note that our proof still goes through (and thus Theorem 5 holds) even if $W_j(\mathbf{s}^{t*})$ and $W_j(\mathbf{s}^{t**})$ are obtained via different values of r_j .

The above theorem proves convergence under very broad conditions, but does not provide much intuition as to what happens along the converging path. Our next result shows that when all voters have the same uncertainty level, convergence is much more structured. Our result extends a similar result in [15] for the finite model, but note that in our model convergence is guaranteed even when subsets of voters move simultaneously.

Conveniently, when all voters has the same uncertainty r , at any state \mathbf{a}^t there is just one agreed set of possible winners, denoted by $W^t = W(\mathbf{a}^t)$. We say that a move $a_i \xrightarrow{i} a'_i$ is an *opportunity move* if $a'_i \succ_i a_i$, and otherwise it is a *compromise move*.

Proposition 6. *Consider any non-atomic Plurality voting game with weak LD moves, where all voters have the same uncertainty level r . If \mathbf{a}^0 is the truthful state, then for all t : (A) $W^{t+1} \subseteq W^t$; (B) the score of the winner is non-decreasing; (C) there are only compromise moves; and (D) any a_i^t is either the most preferred candidate for i in W^t , or it is not in W^t .*

The proposition still holds if the uncertainty level r decreases over time.

5 Regret Minimization

The local dominance approach is appropriate to describe voters who are reluctant to change their vote, unless they know it cannot hurt them, a behavior that is consistent with loss aversion. A different approach to decision making under strict uncertainty is minimization of *worst case regret*, explained by *risk aversion* [3]. Intuitively, such a voter wants to avoid situations in which she could have significantly improved her utility, but failed to be pivotal.

On one hand, regret minimization is simpler than local-dominance since it does not depend on the current vote. Thus we can write $g_v(\mathbf{s}), W_v(\mathbf{s})$ rather than $g_i(\mathbf{s}), W_i(\mathbf{s})$. On the other hand, regret may depend on *cardinal utilities*, rather than ordinal preferences.

A cardinal utility scale is a generic function $u : M \rightarrow \mathbb{R}$, where $u(a) \neq u(b)$ for all $a \neq b$. A cardinal utility scale u fits order $Q \in \pi(M)$, if $u(a) > u(b)$

whenever $Q(a) < Q(b)$. We thus augment the definition of a type v with a cardinal utility scale u_v that fits Q_v .

Formally, the *regret* of a type v voter for voting b in state \mathbf{s}' is $REG_v(\mathbf{s}', b) = \max_{c \in M} u(f(\mathbf{s}', c)) - u(f(\mathbf{s}', b))$. Note that $REG_v(\mathbf{s}', b) \geq 0$.

The *worst case regret* (WCR) of $i \in I$ for voting b in \mathbf{s} is $WCR_v(\mathbf{s}, b) = \max_{\mathbf{s}' \in S(\mathbf{s}, r_v)} REG_v(\mathbf{s}', b)$.

Definition 3 (Strategic move under regret minimization). A WCR voter of type v in profile \mathbf{a} votes for a candidate b minimizing $WCR_v(\mathbf{a}, b)$.

We first characterize WCR moves. We show that (under a candidate-wise distance) we can effectively ignore the cardinal utility scale, as only the preference order affects the vote.

Lemma 7. *Either $|W_v(\mathbf{s})| = 1$ (in which case all regrets are 0); or the unique candidate minimizing $WCR_v(\mathbf{s}, c)$ is $a^* = \operatorname{argmin}_{c \in W_v(\mathbf{s})} Q_i(c)$.*

From Lemmas 4 and 7, we get that regret minimization provides a partial justification for the strict variant of LD moves.

Corollary 8. *Let $a_i \xrightarrow{i} a'_i$ be a strict LD move in state \mathbf{s} , then a'_i is the WCR response of type v_i , i.e., the unique candidate c minimizing $WCR_{r_i}(\mathbf{s}, c)$.*

This does not mean that the WCR dynamics and the strict LD dynamics coincide: in an arbitrary state \mathbf{a} it is possible that a_i is a possible winner that is undominated (and hence i will not move under LD), yet there is another possible winner $a'_i \succ_i a_i$ (and hence i will move to a'_i under WCR).

It is an open question whether a WCR game where all voters have the same r is acyclic, yet if we add another restriction then we get that WCR moves coincide with LD moves, and in particular must converge.

Proposition 9. *Consider any non-atomic Plurality voting game, where all voters have the same uncertainty level r . If \mathbf{a}^0 is the truthful state (and $|W_r(\mathbf{a}^0)| \geq 2$), then for every time t and any $i \in I$, the WCR move and the strict LD move of i coincide. In particular, the WCR dynamics converges to an equilibrium.*

Proof. Consider any agent $i \in I$. By property (D) of Proposition 6, we have that either $a_i^t \notin W_r(\mathbf{s}^t)$, or a_i^t is the best possible winner. In the first case, by Lemma 4, there is a strict LD move, and by Corollary 8, this move coincides with the WCR move.

In the latter case, there is no LD move for i (so $a_i^{t+1} = a_i^t$ under LD), and by Lemma 7, a_i^t minimizes worst case regret (so $a_i^{t+1} = a_i^t$ under WCR).

Finally, since the LD dynamics converges, we get that under the conditions of the proposition (same r , truthful initial state), WCR converges as well. \square

Diverse population For the local-dominance dynamics, Theorem 5 shows that any game is acyclic. Since under regret minimization voters are more likely to have a strategic move, it is also more likely that cycles emerge, and even the existence of an equilibrium is not guaranteed (see Appendix XII for proof).

Proposition 10. *There is a voting game where no voting equilibrium exists under WCR dynamics.*

6 Discussion and Related Work

We refer the reader to [15], which surveyed many game-theoretic models of voting equilibrium, in particular w.r.t. various approaches to uncertainty and dominance. In addition, Meir et al. showed via extensive simulations that the equilibria reached by LD voters reproduce patterns observed in the real world, such as Duverger’s Law [6].

Other concepts similar to local-dominance moves have been recently introduced in the voting literature, notably *II-manipulation* [5, 20] and *’de re’ manipulation* [24]. However these papers did not consider distance metrics and did not present any results on equilibrium existence or convergence.

Voting behavior based on regret minimization was considered by Ferejohn and Fiorina [8]. However their model (like probability-based models) heavily relies on voters having cardinal utilities. Also, they take an extreme approach where voters do not use *any* available information, and thus all states are considered possible. Another regret-based model was suggested in [18], which also ignores any available information. Merrill shows that under the Plurality rule uncertain voters should be *truthful*, which stands in sharp contrast to behavior observed in the real world.

We see our regret minimization model as a non-probabilistic variation of the Myerson and Weber model [19]. Specifically, in the MW model the voter considers the probability of each tie to conclude her expected utility. In our WCR model the voter focuses on the *most significant* possible tie, which greatly facilitates the decision making process.

Voting experiments suggest that human voting behavior is consistent with regret minimization [4, 14], though voters may not see themselves as such. We should note that these studies have limited relevance to our work since they concern the decision *whether to vote* (voter turnout), rather than the strategic decision *what to vote*. Thus more experimentation is required to test the validity of such models.

A related approach is *iterated regret minimization* [9], which was recently applied to voting [7]. This approach assumes that voters know exactly both the preferences and the decision making process of the other voters, whereas both of these assumptions are avoided in the models we studied.

A recent result on memoryless dynamics in atomic games, suggests that in any game with more than one pure equilibrium, cycles must occur if arbitrary

groups of agents can move simultaneously [11]. Our result shows that this is no longer true in nonatomic games, and in fact there is a large class of games where convergence is guaranteed despite the existence of multiple equilibria.

6.1 Conclusion and Future Directions

We showed that in the Plurality voting system, voters who avoid locally-dominated candidates will always converge to an equilibrium, and that this result is robust to the uncertainty levels in the populations, the initial state, and the order in which voters or groups of voters play.

Convergence rate In finite games without uncertainty, it is known that convergence must occur in polynomial time in the number of voters and candidates. The speed of convergence in non-atomic games is not easily defined. For example agents can just move in smaller and smaller masses in an infinite (acyclic) path. However one can ask if there is a bound on the maximal number of times a *single* voter can move. For example under the conditions of Prop. 6, any voter can move at most $m - 1$ times, so we can say that convergence is “fast.” Our general convergence proof does not provide a bound of such sort, but we believe it is an interesting open question for future research. We conjecture that no agent should move more than a polynomial (in m) number of times until convergence occurs.

Justifying best-response Consider the assumption made in [15], that among all candidates dominating her current action, a voter will always select the one that is most preferred. Our paper tackles this assumption in two ways. First, we show that it is not required for convergence, and can thus be relaxed (among the other restrictions we relax). Second, we show that this assumption can be justified on the grounds that it minimizes the worst case regret of the voter.

Experimental validation

Other voting rules While the definition of weak LD voters naturally extends to many other voting rules (as do the definitions in [5, 20, 24]), the action space in most rules contains all permutations of M , and there may be many actions that dominate a particular action. Thus the mere definitions are not very instructive as to how a voter would act. We hope the (local) worst-case regret minimization approach will be useful in defining reasonable voting behaviors under uncertainty for other voting rules. Finally, distance based uncertainty may prove useful in other games where there is a natural metric over action profiles.

Acknowledgments

The author would like to thank Omer Lev, David Parkes and Maria Polukarov for insightful discussions and for commenting on drafts of this paper, and Har-

ward CRCS for the support. Some of the future directions and clarifications were pointed out by anonymous referees.

Online appendices for paper #20

VII Characterization of LD moves

Proposition 6. *Consider any non-atomic Plurality voting game with weak LD moves, where all voters have the same uncertainty level r . If \mathbf{a}^0 is the truthful state, then for all t : (A) $W^{t+1} \subseteq W^t$; (B) the score of the winner is non-decreasing; (C) there are only compromise moves; and (D) any a_i^t is either the most preferred candidate for i in W^t , or it is not in W^t .*

Proof. The proof is very similar to the one from [15]. Since the game is non-atomic the proof is somewhat simpler and also applies to group-moves. In contrast to [15] (who focused on the ℓ_1 metric), our proof also holds for *weak* LD moves under candidate-wise metrics.

We show that by induction on t , all properties (A-D) hold. This clearly holds in $t = 0$, since $a_i^0 = q_i$. Now, consider some move from \mathbf{a}^t to \mathbf{a}^{t+1} . By the induction hypothesis all properties hold until \mathbf{a}^t .

Consider any move $a_i^t \xrightarrow{i} a_i^{t+1}$ for some $i \in I$. If this is the first move of voter i , then clearly this is a compromise move. Otherwise, since a_i^t is the most preferred candidate of i in some $W^{t'}$, $t' < t$, and $a_i^{t+1} \in W^t \subseteq W^{t'}$, then $a_i^t \succ_i a_i^{t+1}$, which shows (C). By Lemma 4, $a_i^{t+1} \in W^t$. Since $a_i^t \succ_i a_i^{t+1}$, categories (b) and (c) of Lemma 3 are impossible, thus category (a) applies, and $a_i^t \notin W^t$. In particular $a_i^t \neq f(\mathbf{a}^t)$, thus neither (A) nor (B) are violated at time t .

Finally, by Lemma 4 a_i^{t+1} is the most preferred in W^t . Since $W^{t+1} \subseteq W^t$, (D) holds as well. \square

The proposition still holds if the uncertainty level r *decreases* over time, but not if it can grow (since then it may occur that $W^{t+1} \not\subseteq W^t$).

VIII Atomic voters

Basic notations We denote $[x] = \{1, 2, \dots, x\}$. The set of candidates is denoted by M , respectively, where $m = |M|$. Let $\pi(M)$ be the set of all orders over M .

There is a finite set of voters I of size n . A *preference profile* $\mathbf{Q} \in \pi(M)^n$ is a collection over preferences, where Q_i is the preference order of voter $i \in I$. We denote by v_i the *type* of voter i , which includes her preferences and possibly other information on her beliefs and behavior. Thus a population $\mathbf{V} = (v_i)_{i \in I}$ specifies the types of all voters. For a vector $\mathbf{x} = (x_j)_{j \in I}$, we denote by \mathbf{x}_{-i} the partial vector that includes all entries except x_i .

We denote $a \succ_i b$ if $Q_i(a) < Q_i(b)$.

An *action profile* (a.k.a. state) is a vector $\mathbf{a} = (a_i)_{i \in I}$, where $a_i \in M$ is the candidate that voter i votes for.

The Plurality voting rule f allows voters to vote for a single candidate from the set M . Then, f chooses the candidate with the highest score. Thus any population defines a *game*, where voters' types determine their utilities. The

outcome of the Plurality rule only depends on the total number of votes for each candidate. Formally, we define the *score vector* $\mathbf{s}_{\mathbf{a}}$ induced by action profile \mathbf{a} , as $s_{\mathbf{a}}(c) = \sum_{i \in I: a_i = c} w_i$, where $w_i = 1$ for all i (compare with the corresponding definition in the nonatomic case).

Every score vector \mathbf{s} implicitly contains an arbitrary “tie-breaker” $Q_{\mathbf{s}}^+ \in \mathcal{Q}$. The *winner* $f(\mathbf{s})$ is the candidate $c \in M$ whose score $s(c)$ is maximal if such a candidate is unique. Otherwise we break ties according to $Q_{\mathbf{s}}^+$, thus $f(\mathbf{s}) = \operatorname{argmin}_{c \in \operatorname{argmax}_{s(c)} Q_{\mathbf{s}}^+(c)}$. An alternative (but equivalent) way to compute the winner in the atomic case, is to think of the tie breaker as if it adds a small fraction to the score of each candidate. Formally, $f(\mathbf{s}) = \operatorname{argmax}_{c \in M} (s(c) + \frac{m+1-Q_{\mathbf{s}}^+(c)}{m+1})$.

Dynamics We describe the behavior of a voter $i \in I$ by a *response function* $g_i : \mathcal{S} \rightarrow 2^M \setminus \emptyset$. That is, a mapping from the current state to a subset of valid actions. In contrast to the atomic case, the state (score vector) considered by voter i does not contain her own action. Thus $g_i(\mathbf{a})$ is a shorthand for $g_i(\mathbf{s}_{\mathbf{a}_{-i}})$, and a voter $i \in I$ may choose any action in $g_i(\mathbf{a})$. Together, Q_i and g_i completely define the type v_i of voter i .

For example, g_i can be the best-response function: $g_i(\mathbf{a}) = \operatorname{argmin}_{c \in M} Q_i(f(\mathbf{a}_{-i}, c))$.

Definition 4. Let \mathbf{V} be a population of voters. A *voting equilibrium* is a state \mathbf{a} , where $g_i(\mathbf{a}) = \{a_i\}$ for all $i \in I$.

For example, if g_i is the best-response function (or better-response, which is set-valued), then a voting equilibria coincides with the Nash equilibria of the atomic game.

Local dominance The set $S_i(\mathbf{s}, x)$ is exactly as in the nonatomic case, i.e.

$$S(\mathbf{s}, x) = \{\mathbf{s}' : \delta(\mathbf{s}, \mathbf{s}') \leq x\},$$

however for a profile \mathbf{a} , $S_i(\mathbf{a}, x) = S_i(\mathbf{s}_{\mathbf{a}_{-i}}, x)$. In particular, and in contrast to the nonatomic case, two voters of the same type may consider different sets of possible states, if they happen to vote for different candidates in \mathbf{a} . For the multiplicative metric, $s'(c)$ is in the range $[\lfloor s(c)/(1+x) \rfloor, \lceil s(c)(1+x) \rceil]$.

We define the *uncertainty parameter* $r_i \in \mathbb{R}^+$ as part of the type v_i . Given a prospective state \mathbf{s} , voter i believes that the actual state may be any $\mathbf{s}' \in S_i(\mathbf{a}, r_i)$.

Since \mathbf{s}' also does not contain the vote of i , the winner that i considers in \mathbf{s}' depends on her vote a_i . That is, $f(\mathbf{s}', a_i)$ is the winner of the profile \mathbf{s}' where a_i gets one additional vote. Note that in the nonatomic case, the vote a_i is assumed to break any ties, and the tie breaker $Q_{\mathbf{s}'}^+$ is only considered if a_i is not a candidate with maximal score. In contrast, in the atomic case a_i simply adds one vote to $s(a_i)$, and then we apply $Q_{\mathbf{s}}^+$ to resolve any ties. Thus if several candidates $C \subseteq M$ are tied in (\mathbf{s}, a_i) with maximal score, each of them wins in some possible state.

Definition 5 (Strategic move under local dominance for finite populations). Let a_i be the vote of i in profile \mathbf{a} . Let $D \subseteq M$ be the set of candidates that $S(\mathbf{a}, r_i)$ -dominate a_i , and that are not $S(\mathbf{a}, r_i)$ -dominated. If $D = \emptyset$ then $g_i(\mathbf{a}) = \{a_i\}$. Otherwise:

- For a *weak LD voter*, $g_i(\mathbf{a}) = D$.
- For a *strict LD voter*, $g_i(\mathbf{a}) = \operatorname{argmin}_{d \in D} Q_i(d)$.

In other words, a strict LD voter votes the most preferred candidate that locally-dominates her current choice (if such exists). The definition of a weak LD voter is much more permissive, and does not restrict the voter to select the most preferred candidate in D .

Note that the definition looks identical to Def. 2 from the nonatomic model. The only difference is that $S_i(\mathbf{a}, r_i)$ in the finite case depends on the identity (on the current vote) of i , and not only on its type. Note however that g_i depends on a_i both in the finite and nonatomic cases.

Definition 5 for strict local dominance almost coincides with the definition in of local dominance in [15], but for one difference: We require D to only contain undominated candidates whereas no such restriction appeared in [15]. Note that without this restriction, a single voter can move, and then immediately move again, which does not make much sense.

Characterization of LD moves We say that candidate c is a *possible winner* for i in state \mathbf{a} if there is a possible state where c wins. Formally, $W_i(\mathbf{a}) = \{c \in M : \exists \mathbf{s}' \in S_i(\mathbf{a}, r_i) \text{ s.t. } f(\mathbf{s}', c) = c\}$.

Note that in contrast to the nonatomic case, different voters of the same type may consider different sets of possible winners. This is the main advantage of the nonatomic model.

It is easy to see that under the multiplicative distance $c \in W_i(\mathbf{a})$ iff $s_{\mathbf{a}-i}(c) \geq (1 + r_i)^{-2} s_{\mathbf{a}-i}(f(\mathbf{s}))$. The full proof of Lemma 2 (in Appendix XI) that also covers the finite voters case, for any candidate-wise metric.

The proofs of Lemmas 3 and 4 are the same for the finite case, one we replace the notation $S(\mathbf{a}, r_i)$ with $S_i(\mathbf{a}, r_i)$.

Convergence Recall that when several voters (in the finite case) may move simultaneously, cycles may emerge even without uncertainty at all [17]. A simple example is when $s(a) = 1, s(b) = s(c) = 0$, and we add two voters that favor b, c to a . If they start by one voting to b and the other to c , then each time they both play they will just switch locations.

The proof of Theorem 5 fails in the finite case where it uses the fact that $s^{t*}(a') \geq s^{t*}(a^*) = s^*$ and thus the deserted candidate a_j is a possible winner for j . In the finite case this does not hold, since the high score of a_j depends on the vote of j , and j herself does not consider a_j as a possible winner.

We next show that the theorem still holds once restricted to singleton moves. That is, we consider all valid sequences $\mathbf{a}^0, \mathbf{a}^1, \dots$ where $a_i^{t+1} \in g_i(\mathbf{a}^t)$ for some voter $i \in I$, and $a_i^{t+1} = a_i^t$ for all other voters.

As we can see, in the finite case we need to cover some more issues so the proof is somewhat longer, despite the restriction to singleton moves.

Theorem 11. *Any sequence of (singleton) weak LD moves is finite.*

Proof. Clearly in the finite case every sequence is either finite or contains a cycle. Assume, toward a contradiction, that there is a cyclic path $(\mathbf{a}^t)_{t=0}^T$. Let $R \subseteq M$ be all candidates that are part of the cycle, and let $I(R) \subseteq I$ be the voters that participate in the cycle.

Let s^* be the lowest score of any candidate in R during the cycle. That is, we find $a^* \in R$ and t^* s.t. $s^* = s^{t^*}(a^*)$ is minimal.

By definition, $s^* = s^{t^*}(a^*) \leq s^t(c)$ for all $c \in R$ for every time $t \leq T$. Consider the next step where some voter j joins a^* (w.l.o.g. at step t^*), and pick an arbitrary voter $j \in I$ s.t. $a^* = a_j^{t^*+1} \neq a_j^{t^*} = a_j$ (a voter that moved to a^*). Thus at step t^* there is a move $a_j \xrightarrow{j} a^*$ for some $a_j \in R$.

By Lemma 4, a^* is a possible winner for voter j , i.e., $a^* \in W_j(\mathbf{s}^{t^*})$. We facilitate the notation by writing $\mathbf{s}_{-j} = \mathbf{s}_{\mathbf{a}_{-j}}$, which are the scores from the point of view of j (i.e., omitting the vote of j).

Since $s^{t^*}(a') \geq s^*$ for all $a' \in R$, in particular $s_{-j}^{t^*}(a_j) + 1 \geq s^*$. Suppose first that $s_{-j}^{t^*}(a_j) < s^*$. Then after the move $s^{t^*+1}(a_j) = s_{-j}^{t^*}(a_j) < s^*$, in contradiction to the selection of s^* (note that this only applies to singleton moves, as no other voter can move to a_j at time t^*). Thus $s_{-j}^{t^*}(a_j) \geq s^* = s_{-j}^{t^*}(a^*)$, which means that $a_j \in W_j(\mathbf{a}^{t^*})$. Note that for all other candidates $c \in R \setminus \{a_j\}$, $s_{-j}^{t^*}(c) = s^{t^*}(c) \geq s^*$ as well, thus $R \subseteq W_j(\mathbf{s}^{t^*})$. The remainder of the proof is almost identical to the nonatomic case, but we write it for completeness.

By Lemma 3, either a_j is the least-preferred candidate for j in $W_j(\mathbf{s}^{t^*})$ (Cases I and II), or the third category of the lemma holds. We treat the latter case separately (Case III), so assume a_j is indeed the least-preferred in $W_j(\mathbf{s}^{t^*})$. Since $R \subseteq W_j(\mathbf{a}^{t^*})$, a_j is the least-preferred in R as well.

There must be some step t^{**} in the cycle where a voter of type v_j moves to a_j (w.l.o.g. voter j). So in $\mathbf{s}^{t^{**}}$, a_j is preferred by j to *some* other possible winner z by Lemma 4. Since a_j is the least-preferred in R , and $a_j \succ_j z$, $z \notin R$. Thus there is some $z \in W_j(\mathbf{s}^{t^{**}}) \setminus R$ such that $z \prec_j a_j$. Denote the (fixed) score of z by $s(z)$.

Case I: $s(z) \geq s^*$. Consider again step t^* . Since $s(z) \geq s^* = s_{-j}(a^*)$ and $a^* \in W_j(\mathbf{a}^{t^*})$, we have $z \in W_j(\mathbf{a}^{t^*})$. Since a_j is the least preferred possible winner in t^* , we have that $z \succ_j a_j$, which is a contradiction.

Case II: $s(z) < s^*$. Denote $d = a_j^{t^{**}}$, and consider the step $d \xrightarrow{j} a_j$ at time t^{**} . Since $d \in R$ then $s_{-j}^{t^{**}}(d) \geq s^* > s(z)$, and thus $d \in W_j(\mathbf{a}^{t^{**}})$ (as $z \in W_j(\mathbf{a}^{t^{**}})$). By Lemma 3 (category (b) or (c)), we have that $a_j \succ_j d$. This is a contradiction since a_j is the least-preferred in R .

Case III: The remaining case is when there is $b \in W_j(\mathbf{a}^{t^*})$ s.t. $a_j \succ_j b$. Then by category (c) of Lemma 3, $a^* \in W_0(\mathbf{a}^{t^*})$. However since a^* has minimal score, this means that $R \subseteq W_0(\mathbf{a}^{t^*})$, i.e., all candidates $c \in R$ have the same score $s_{-j}^{t^*}(c) = s^*$ at time t^* , and there is exactly one candidate (a_j) that has the

additional vote of voter j . Since throughout the cycle the score of candidates in R may not drop below s^* , we have that all moves in the cycle are from the single candidate $c^t \in R$ with score $s^t(c^t) = s^* + 1$ to a candidate with score s^* . This means that the deserted candidate c^t is always a possible winner for the mover. Consequently, all of the moves in the cycle fall under categories (b) or (c) of Lemma 3. Thus voters only vote for more preferred candidates, which contradicts a cycle. \square

Best response and better response In [17] a distinction was made between “better-response” (any move that improves the winner), “best-response” (any move that leads to the selection of the best possible winner); and “unique best-response” (or restricted best-response), which requires in addition that the voter votes for the selected candidate. It was shown that only unique best-response guarantees convergence, whereas cycles may occur under better-response. It may thus seem that our result contradicts the one in [17].

However, a similar distinction can be made between “better-response” and “restricted better-response”. In [16] it was shown that local-dominance naturally restricts the response in a similar way—a candidate that is not a possible winner is always dominated. In the special case of $r = 0$ the set D is exactly the set of restricted better-responses. Thus it is not the restriction to best response that leads to convergence, but rather the restriction to vote for the candidate the voter is trying to promote.

The connection between local-dominance and restricted better-response for $r = 0$ only holds in our neutral tie-breaking, but it is easy to verify that restricted better-response converges also in the lexicographic tie-breaking model of [17].

IX Convergence and Acyclicity

Without restricting the resolution ϵ , there may be infinite *acyclic* paths, even with two candidates and one type of voters. For example consider a population where all voters prefer b over a , but initially all vote to a . That is, there is just one type of voters and $s^0(a) = 1, s^0(b) = 0$. Suppose that in step t , half of the remaining a voters are chosen to play, and move to b . This entails an infinite acyclic path.

To avoid such infinite paths, and also to allow for more intuitive notation, we assume that there is some bounded granularity level ϵ , and that sets of voters of the same type with mass ϵ always move together.⁵

Thus a state \mathbf{a} is ϵ -valid if $a(v, c)$ is an integer multiple of ϵ , and a move from state \mathbf{a}^t to state \mathbf{a}^{t+1} is ϵ -valid, if (I): both states are ϵ -valid; (II) \mathbf{a}^{t+1} is attained from \mathbf{a}^t by taking each set $i \in I$, and set either $a_i^{t+1} = a_i$, or $a_i^{t+1} \in g_i(\mathbf{a})$.

Finally, for any $\epsilon > 0$, there is only a finite number of ϵ -valid states. Thus there cannot be cycles of infinite length.

⁵Note that this is different than a coalitional manipulation, where voters are coordinated. It is possible that all voters will be less satisfied after the move.

Observation 12. *If a voting game has no finite cycles, then every ϵ -valid path, for any $\epsilon > 0$, is finite.*

Finally, the assumption that all voter sets have the same size ϵ simplifies the notation but is not necessary.

Clearly valid moves can be defined without referring to “identified” voters, although the definition is quite tedious. Consider a valid finite cycle of states (not necessarily ϵ -valid for any ϵ). Since there is a finite number of types and a finite number of states in the cycle, we can still partition the voters $[0, 1]$ to subsets I , s.t. all voters in set $i \in I$ have the same type v_i and the same sequence of actions a_i^1, \dots, a_i^T . Thus we get identified voters in retrospect. The mass of each such set $i \in I$ may not be rational so there may not be common unit size ϵ , but our convergence proof works just the same.

X Voter influence

In the atomic case the meaning of $f(\mathbf{s}, c)$ is clear—it is the outcome of adding one more vote to candidate c in state \mathbf{s} . However changing the vote of a “single” non-atomic voter does not change the score. We remedy this by assuming each voter considers her own weight as *tends to zero* but not quite zero. Formally, denote by $\mathbf{s}_\epsilon^{\rightarrow c}$ the state that is attained from \mathbf{s} if we add to \mathbf{s} a mass of $\epsilon > 0$ voters voting for c . That is, $s_\epsilon^{\rightarrow c}(c) = s(c) + \epsilon$, and $s_\epsilon^{\rightarrow c}(b) = s(b)$ for any other candidate.

Finally, we define (in the non-atomic case) the modified outcome $f(\mathbf{s}, c)$ as $\lim_{\epsilon \rightarrow 0} f(\mathbf{s}_\epsilon^{\rightarrow c})$. To see that the limit is well defined, note that it is sufficient to consider any $0 < \epsilon < \min_{b, c \in M} |s(b) - s(c)|$.

By continuity of the distance metric, if there are states where two different candidates win, then there is also a state where they are tied. Thus voter i believes she can affect the outcome if only if she can act as a tie-breaker between candidates in some possible state. More formally, suppose that voter i votes i in state \mathbf{s} , and consider $S = S(\mathbf{s}, r_i)$.

Lemma 13. *1. if $f(\mathbf{s}') = c$ for all $\mathbf{s}' \in S$, then $f(\mathbf{s}', a_i) = c$ (that is, i cannot change the outcome).*

2. if several candidates $C \subseteq M$ are tied (with maximal score) in some \mathbf{s}' , and $a_i \in C$, then $f(\mathbf{s}', a_i) = a_i$.

3. if several candidates $C \subseteq M$ are tied (with maximal score) in some \mathbf{s}' , and $a_i \notin C$, then any candidate in C wins in some state $\mathbf{s}'' \in S$.

Proof. 1. For all $\mathbf{s}' \in S$, $s'(c) > s'(a_i) + \min_{b, c \in M} |s(b) - s(c)| > s'(a_i) + \epsilon$ for some $\epsilon > 0$. Thus $f(\mathbf{s}', a_i) = \lim_{\epsilon \rightarrow 0} f(\mathbf{s}'_\epsilon^{\rightarrow a_i}) = c$.

2. For all $c \in M \setminus \{a_i\}$, $s'(a_i) \geq s'(c)$ and thus $s_\epsilon^{\rightarrow a_i}(a_i) > s'(a_i) \geq s'(c) = s_\epsilon^{\rightarrow a_i}(c)$.

3. For $c \in C$, consider the state \mathbf{s}'' where $Q_{\mathbf{s}''}^+(c) = 1$. Then the score of a_i in $\mathbf{s}_\epsilon^{\rightarrow a_i}$ is strictly lower than the score of c , as in 1. Thus C have maximal score in $\mathbf{s}_\epsilon^{\rightarrow a_i}$, and c wins due to the tie-breaking advantage. \square

XI Other Distance Functions

Candidate-wise distances

Lemma 2. *Every pair of possible winners are tied for victory in some possible state. Formally, for every $b, c \in W_i(\mathbf{s})$, there is $\mathbf{s}' \in S(\mathbf{s}, r_i)$ s.t. $b, c \in W_0(\mathbf{s}')$.*

Proof for any candidate-wise distance. In this proof we denote the score of c by $s(c)$. Recall that $\delta(\mathbf{s}, \mathbf{s}') = \max_c \delta'(s(c), s'(c))$.

Consider some $b, c \in W_i(\mathbf{s})$. Let the score of all candidates except b, c be $s'(a) = \min\{z : \forall a \notin \{b, c\}, \delta'(s(a), z) \leq r_i\}$. Let $h(b) = \max\{z : \delta'(s(b), z) \leq r_i\}$ and $h(c) = \max\{z : \delta'(s(c), z) \leq r_i\}$, then set $s'(b) = s'(c) = \min\{h(b), h(c)\}$. Then $\mathbf{s}' \in S(\mathbf{s}, r_i)$ by definition. Also, we argue that b, c are tied with maximal score in \mathbf{s}' . Indeed, since $b \in W_i(\mathbf{s})$, we know that there is some state $\mathbf{s}_b \in S(\mathbf{s}, r_i)$ s.t. $f(\mathbf{s}_b) = b$. We know that $s_b(b) \leq h(b)$, and $s_b(a) \geq s'(a)$ for all other candidates. Since the same holds for c w.r.t. some state \mathbf{s}_c , we get that in \mathbf{s}' either b or c have maximal score. However by construction $s'(b) = s'(c)$ so they are tied.

Note that the proof also covers the finite case. \square

As this is our only proof that makes use of the properties of the distance function, all of our results hold for any candidate-wise distance.

Lemma 4. *Consider an LD move $a_i \xrightarrow{i} a'_i$. Then*

1. $a'_i \in W_i(\mathbf{s})$, and there is some $c \in W_i(\mathbf{s})$ s.t. $a'_i \succ_i c$.
2. For a strict LD move, $a'_i = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$.
3. If $a_i \notin W_i(\mathbf{s})$, then $|g_i(\mathbf{s})| = 1$ (thus weak and strict LD coincide).

Proof. Consider $a_i^* = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$. Clearly a_i^* dominates any candidate not in $W_i(\mathbf{a})$, thus $a'_i \in W_i(\mathbf{s})$.

If there is only one possible winner, no action dominates any other action, and there are no moves. The fact that $a'_i \neq a_i$ entails that there is some possible state where a'_i is tied with some other possible winner c . If a'_i is the least preferred possible winner, then it is dominated by a_i^* , which again means there would be no move. This completes part 1.

Next for part 2. Consider the set of candidates D in Def. 2. Since $a'_i \in D$, it is non-empty. We get from part 1 that $D \subseteq W_i(\mathbf{a})$. It is left to show that a_i^* is in D , as this implies $a_i^* = a'_i$.

Indeed, assume that $a'_i \neq a_i^*$. By Lemma 2 and part 1, there is a state $\mathbf{s}' \in S(\mathbf{a}, r_i)$ where a_i^* is tied for victory with some other possible winner c , and $Q_{\mathbf{s}'}^+(c) = 1$. Thus $f(\mathbf{s}', a_i^*) = a_i^* \succ_i c = f(\mathbf{s}', a_i)$, which means a_i^* $S(\mathbf{s}, r_i)$ -beats a_i .

We consider the three categories of Lemma 3. (a) If $a_i \notin W_i(\mathbf{s})$, then in every $\mathbf{s}' \in S(\mathbf{s}, r_i)$ voting for a_i^* can only improve the outcome from $f(\mathbf{s}') \in W_i(\mathbf{a})$ to a_i^* ; (b) If a_i is the worst possible winner, then a_i does not beat any other

candidate; (c) There is no uncertainty except on the tie-breaker (i.e. $s'(c) = s(c)$ for all c , but $Q_{s'}^+$ may differ), and $a_i^*, a_i \in W_0(\mathbf{s})$. Thus in any possible state $f(\mathbf{s}', a_i^*) = a_i^* \succ_i a_i = f(\mathbf{s}', a_i)$. In all cases a_i does not $S(\mathbf{s}, r_i)$ -beat a_i^* , and thus $a_i^* \in D$.

For the third part, if $|W_i(\mathbf{s})| = 1$, then i is non-pivotal and no action dominates any other action. Thus $g_i(\mathbf{s}) = \{a_i\}$ (no move). Otherwise, consider some $b \in W_i(\mathbf{a})$ that is not a_i^* . Then by Lemma 2 there is a possible state \mathbf{s}' where $b, a^* \in W_0(\mathbf{s}')$ with $Q_{s'}^+(a_i^*) = 1$. Then $f(\mathbf{s}', b) = b \prec_i a_i^* = f(\mathbf{s}', a_i)$, so a_i beats b , which means b cannot dominate a_i , and thus $b \notin g_i(\mathbf{s})$. \square

Non-candidate-wise distance functions

For general distance-based relations, the set of possible winners alone is insufficient to determine which pairs of candidates can be tied. That is, Lemma 2 is violated, and thus our other proofs do not hold as well.

Under the ℓ_1 distance, $\delta(\mathbf{s}, \mathbf{s}') = \sum_{c \in M} |s(c) - s'(c)|$. Consider for example the ℓ_1 metric, $r_i = 5$, and four candidates $\{a, b, c, d\}$ whose scores are $\mathbf{s}_1 = (10, 9, 7, 6)$. While all candidates are possible winners, the last two cannot be tied for victory. However if scores are $\mathbf{s}_2 = (10, 6, 6, 6)$ then every pair can be tied.

It may still hold that voting games converge for other distance functions, and we leave this as an open question for future research.

The example above can also demonstrate that cardinal utilities matter under the WCR dynamics. Suppose that the utilities for the four candidates are $\mathbf{u}_i = (3, 4, 5, 0)$ (i.e. the third candidate is best, the fourth is worst). We can compute the WCR of every candidate in by looking on all pairs that can be tied for victory. In \mathbf{s}_2 we have $WCR_i(\mathbf{s}, c) = \max\{4 - 0, 3 - 0\} = 4$, whereas the WCR of any other action is 5. Thus c minimizes WCR. However in \mathbf{s}_1 , the tie between c, d is impossible. Thus $WCR(\mathbf{s}_1, b) = \max\{3 - 0, 5 - 3\} = 3$, whereas $WCR(\mathbf{s}_1, c) = WCR(\mathbf{s}_2, c) = 4$. Note that in this case the voter will not vote for her most favorite candidate (c), even though it is a possible winner. Finally, note that if we change $u(d)$ to 2.5 (i.e., without changing ordinal preferences Q_i), then c is once again selected in \mathbf{s}_1 under WCR.

Thus whether voters are sensitive to cardinal utilities, strongly depends on the uncertainty model we are assuming, and in particular on the distance function.

XII Regret Minimization

Lemma 7. *Either $|W_v(\mathbf{s})| = 1$ (in which case all regrets are 0); or the unique candidate minimizing $WCR_v(\mathbf{s}, c)$ is $a^* = \operatorname{argmin}_{c \in W_v(\mathbf{s})} Q_i(c)$.*

Proof. If there is a unique possible winner d then the $f(\mathbf{s}', b) = d$ for all $\mathbf{s}' \in S(\mathbf{s}, r_v)$, $b \in M$, and thus $WCR_v(\mathbf{s}, b) = 0$. Thus suppose that $a^*, b \in W_v(\mathbf{s})$, where b is the least-preferred for type v .

We classify possible states into two classes S_1, S_2 . If $f(\mathbf{s}', a^*) = a^*$ we say that $\mathbf{s}' \in S_1$, and otherwise $\mathbf{s}' \in S_2$.

In $\mathbf{s}' \in S_1$, $REG_v(\mathbf{s}', a^*) = 0 \leq REG_v(\mathbf{s}', c)$ for all $c \in M$. Recall that by Lemma 2, any two possible winners can be tied for victory. Thus in some $\mathbf{s}' \in S_1$, $a^*, b \in W_0(\mathbf{s}')$ and $Q_{\mathbf{s}'}^+(b) = 1$. Then for any $c \neq a^*$, $REG_v(\mathbf{s}', c) \geq u(f(\mathbf{s}', a^*)) - u(f(\mathbf{s}', c)) = u(a^*) - u(b)$.

In $\mathbf{s}' \in S_2$, $f(\mathbf{s}', a^*) = d \neq a^*$, then $d \in W_0(\mathbf{s}')$, $Q_{\mathbf{s}'}^+(d) = 1$, and $REG_v(\mathbf{s}', a^*) = \max_{c' \in W_0(\mathbf{s}')} u(c') - u(d)$, since a voter can only bring about the election of a candidate that is tied for victory. Thus $REG_v(\mathbf{s}', a^*) \leq |u(c') - u(b)|$ for some $c', b \in W_0(\mathbf{s}')$, $c' \neq a^*$.

Thus $WCR_v(\mathbf{s}, a^*) \leq |u(c') - u(b)|$ for some $c' \in W_v(\mathbf{s})$, $c' \neq a^*$, which means that for all $c \in M$, $WCR_v(\mathbf{s}, a^*) < u(a^*) - u(b) \leq WCR_v(\mathbf{s}, c)$. \square

In [19], it is shown that a voter can calculate the expected utility of voting for c , by averaging $u_i(c) - u_i(b)$ over all candidates b , weighted by the probability that b, c are tied. Lemma 7 demonstrates that under the regret minimization scheme, voters apply similar logic (only consider possible ties). However instead of taking a weighted average, the voter only considers the maximum.

Proposition 10. *There is an example of a non-atomic voting game, where no voting equilibrium exists under WCR dynamics.*

Proof. There are 6 candidates, $\{a, b, c_1, c_2, c_3, d\}$. We set the base scores of candidates to $\bar{\mathbf{s}} = (12, 6, 0, 0, 0, 12)$. That is, scores due to non-strategic voters. We use the ℓ_∞ norm.

We introduce four types of strategic voters: v and $(w_j)_{j=1}^3$. The preferences are $Q_v : b \succ a \succ c_i \succ d$ for all i , and $Q_{w_j} : c_j \succ b \succ a \succ d \succ c_i$ for $i \neq j$. We also set $r_v = 2, r_w = 6$. There are 1.5 units of type v voters, and 1.2 units of each type w_j voters. We note that under the ℓ_∞ metric, $W_r(\mathbf{s})$ contains all candidates whose score is at least $s(f(\mathbf{s})) - 2r$. Also note that the total mass of all voters is 5.1.

We show that in any profile \mathbf{a} there is some WCR move. We denote by x_e , $e \in M$ the fraction of voters (of any type) voting for candidate e in \mathbf{a} . We observe that $s(a) \leq 12 + 5.1 < 18 \leq s(b) + 2r_w$, thus b is always a possible winner for type w_j voters.

1. If there are type w_j voters on d, a or c_i ($i \neq j$), they prefer b , since $b \in W_w(\mathbf{a})$ and $b \succ_w d, a, c_i$.
2. Similarly, if there are type v voters on d or any c_i , they prefer a (w.l.o.g. there are no type w agents on d , so $a \in W_v(\mathbf{a})$).
3. A type v voter votes for b if $b \in W_v(\mathbf{a})$ and otherwise for a .
4. A type w_j voter votes for c_j if $c_j \in W_w(\mathbf{a})$ and otherwise for b .

Assume toward a contradiction that \mathbf{a} is an equilibrium under WCR. We know by 1 and 2 that $x_d = 0$, $x_{c_j} \leq 1.2$ for all j . We observe that $b \in W_v(\mathbf{a})$ iff the inequality in $6 + x_b + r_v = s(b) + r_v \geq s(a) - r_v = 12 + x_a - r_v$ holds. That is, iff

$x_b \geq x_a + 12 - 6 - 4 = x_a + 2$. Similarly, $c_j \in W_w(\mathbf{a})$ iff $s(c_j) + r_w \geq s(a) - r_w$, which means iff $x_{c_j} \geq x_a$.

Suppose first that $b \in W_v(\mathbf{a})$. Then all type v voters are on b , which means $x_a = 0, x_b \geq 1.5$. This entails that $c_j \in W_w(\mathbf{a})$ for all j , since $x_{c_j} \geq 0 = x_a$. We conclude that all w_j voters are on c_j , which means $x_b \leq 1.5$. However we get that

$$s(b) = \bar{s}(b) + x_b \leq 6 + 1.5 = 7.5 < 8 = s(a) - 2r_v,$$

so we get that $b \notin W_v(\mathbf{a})$. A contradiction.

Thus suppose that $b \notin W_v(\mathbf{a})$. Then all type v voters are on a , which means $x_a = 1.5 > x_{c_j}$ for all j . This entails that $c_j \notin W_w(\mathbf{a})$ for all j , since $x_a = 1.5 > 1.2 \geq x_{c_j}$. We conclude that all w_j voters are on b , which means $x_b \geq 3.6$. However we get that

$$s(b) = \bar{s}(b) + x_b \geq 6 + 3.6 = 9.6 > 9.5 = s(a) - 2r_v,$$

so we get that $b \in W_v(\mathbf{a})$. A contradiction again. \square

Proposition 14. *There is an example of a atomic population \mathbf{V} with four candidates, where no voting equilibrium exists under WCR dynamics.*

Proof. We set the base scores of candidates to $\bar{\mathbf{s}} = (9, 4, 0, 10)$. That is, scores due to non-strategic voters. We use the ℓ_∞ norm.

We introduce two strategic voters i and j . The preferences are $Q_i : b \succ a \succ c \succ d$, $Q_j : c \succ b \succ a \succ d$, with lexicographic tie-breaking over candidates. We set the uncertainty levels to $r_i = 1$, $r_j = 4$. Thus a candidate is in W_i if it needs more 3 votes (including i) to win, and in W_j if it needs more 9 votes.

We first show that there is a cycle

$$\mathbf{a}^1 = (b, c) \xrightarrow{i} \mathbf{a}^2 = (a, c) \xrightarrow{j} \mathbf{a}^3 = (a, b) \xrightarrow{i} \mathbf{a}^4 = (b, b) \xrightarrow{j} (b, c) = \mathbf{a}^1.$$

We can write the same cycle in terms of score vectors:

$$\mathbf{s}^1 = (\mathbf{9}, 5, 1, \mathbf{9}) \xrightarrow{i} \mathbf{s}^2 = (\mathbf{10}, \mathbf{4}, 1, \mathbf{9}) \xrightarrow{j} \mathbf{s}^3 = (\mathbf{10}, 5, 0, \mathbf{9}) \xrightarrow{i} \mathbf{s}^4 = (\mathbf{9}, \mathbf{6}, \mathbf{0}, \mathbf{9}) \xrightarrow{j} (9, 5, 1, 9) = \mathbf{s}^1.$$

We justify the moves as follows by Lemma 7 and the following observations for each move: (1) $b \notin W_i(\mathbf{s}^1)$ but i is pivotal between a, d ; (2) $c \notin W_j(\mathbf{s}^2)$ since $s(a) = 6$, but j is pivotal between b, a ; (3) $b \in W_i(\mathbf{s}^3)$; (4) $c \in W_j(\mathbf{s}^4)$ since $s(a) = 5$.

Note that the third move by agent i cannot be an LD move, since b does not dominate a .

Also, all other profiles are never stable: d is the last candidate of both voters and thus has maximal regret in every state, so no voter votes for d . Voter i never votes for c since it is never a possible winner. Voter j never votes for a since b is always a possible winner and $b \succ_j a$. \square

XIII Problematic Proof in [15]

Proposition 5.5 in [15] states that from any initial state in voting games with a finite population, there is some path that converges to an equilibrium. More specifically, that if opportunity (moves to a more preferred candidate) moves always precede compromise moves then convergence is guaranteed. The Proposition extends this claim to subsets of agents that move simultaneously, provided that eventually enough singleton steps occur. The proof is omitted from [15], but appears in the full version [16].

Ignoring the complications of simultaneous moves, the proof shows that after the first time there are no opportunity moves, then the score of the winner must monotonically increase until convergence. This relies on the observation from a previous proof (about convergence from the truthful state) that only voters who vote for candidates who cannot win may move. However this observation is no longer true with arbitrary initial state.

Consider the following example. Four candidates $\{a, b, c, d\}$ with fixed scores $(5, 3, 4, 1)$. We add two voters with $r = 2$ (we use ℓ_1 metric, but similar examples can be constructed for any metric). i prefers $c \succ a \succ b \succ d$. j prefers $d \succ a \succ b \succ c$. Initially i votes for a and j votes for d .

There are no opportunity moves since c does not dominate a for i (if i moves, b may win with additional two votes). Thus the first step is by j who deserts d and goes to a . Now b is no longer a possible winner, and a is dominated (for i) by c . Thus in the next step i does an opportunity move and the score of the winner decreases.

References

- [1] Robert J Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1):6–19, 1995.
- [2] Robert J Aumann. Interactive epistemology i: knowledge. *International Journal of Game Theory*, 28(3):263–300, 1999.
- [3] David E Bell. Regret in decision making under uncertainty. *Operations research*, 30(5):961–981, 1982.
- [4] André Blais, Robert Young, Christopher Fleury, and Miriam Lapp. Do people vote on the basis of minimax regret? *Political Research Quarterly*, 48(4):827–836, 1995.
- [5] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *AAAI’11*, pages 638–643, 2011.
- [6] Maurice Duverger. *Political Parties: Their Organization and Activity in the Modern State*. New York: John Wiley. Y, 1954.
- [7] Miranda Emery and Mark C Wilson. Iterated regret minimization in voting games. In *COMSOC’14*, 2014.
- [8] John A Ferejohn and Morris P Fiorina. The paradox of not voting: A decision theoretic analysis. *The American political science review*, pages 525–536, 1974.
- [9] Joseph Y Halpern and Rafael Pass. Iterated regret minimization: A new solution concept. *Games and Economic Behavior*, 74(1):184–207, 2012.
- [10] Nathanael Hyafil and Craig Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *UAI’04*, pages 268–277, 2004.
- [11] Aaron D Jaggar, Michael Schapira, and Rebecca N Wright. Distributed computing with adaptive heuristics. In *ITCS’11*, 2011.
- [12] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives*, pages 193–206, 1991.
- [13] Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. In *The Concept of Probability in Psychological Experiments*, pages 25–48. Springer, 1974.
- [14] Joachim I Krueger and Melissa Acevedo. A game-theoretic view of voting. *Journal of Social Issues*, 64(3):467–485, 2008.
- [15] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *ACM-EC’14*, pages 313–330, 2014.

- [16] Reshef Meir, Omer Lev, and Jeffrey S Rosenschein. A local-dominance theory of voting equilibria (full version). *arXiv preprint arXiv:1404.4688*, 2014.
- [17] Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. Convergence to equilibria in plurality voting. In *AAAI'10*, pages 823–828, 2010.
- [18] S Merrill. Strategic voting in multicandidate elections under uncertainty and under risk. In *Power, voting, and voting power*, pages 179–187. Springer, 1982.
- [19] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.
- [20] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *AAMAS'12*, pages 635–644, 2012.
- [21] Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.
- [22] David Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.
- [23] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131, 1974.
- [24] Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *AAMAS'12*, pages 1247–1248, 2012.